

UNIVERSITY OF CALIFORNIA, BERKELEY

BERKELEY • DAVIS • IRVINE • LOS ANGELES • RIVERSIDE • SAN DIEGO • SAN FRANCISCO



SANTA BARBARA • SANTA CRUZ

COLLEGE OF ENGINEERING
DEPARTMENT OF ELECTRICAL ENGINEERING
AND COMPUTER SCIENCES
COMPUTER SCIENCE DIVISION

BERKELEY, CALIFORNIA 94720

August 11, 1995

Dr. Clifford Lau
Electronics Division, ONR312
Office of Naval Research
Scientific Office Code 1114SE
800 North Quincy Street
Arlington, VA 22217-5660


TECHNICAL PROGRESS REPORT

Reference: ONR URI Grant No.: 144-92-J-1617

Dear Dr. Lau:

Enclosed you will find the quarterly progress report for the research project for Professor John Wawrzynek: "Construction of a Connectionist Network Supercomputer," ONR URI Grant No.: N00014-92-J-1617. This report covers the research done during the performance period 5/1/95-7/31/95. Please direct any mail or questions regarding this contract to me.

Thank you for your continued support of this project.



Sincerely,

A handwritten signature in dark ink, appearing to read 'Theresa Lessard-Smith'.

Theresa Lessard-Smith
Project Coordinator

cc: Administrative Grants Officer, Seattle
Director, Naval Research Laboratory
Defense Technical Information Center
Sponsored Projects Office - UC Berkeley
file

19970717 061

DTIC QUALITY INSPECTED 5



DEPARTMENT OF THE NAVY
OFFICE OF NAVAL RESEARCH
SEATTLE REGIONAL OFFICE
1107 NE 45TH STREET, SUITE 350
SEATTLE WA 98105-4631

IN REPLY REFER TO:


4330
ONR 247
11 Jul 97

From: Director, Office of Naval Research, Seattle Regional Office, 1107 NE 45th St., Suite 350,
Seattle, WA 98105
To: Defense Technical Center, Attn: P. Mawby, 8725 John J. Kingman Rd., Suite 0944,
Ft. Belvoir, VA 22060-6218

Subj: RETURNED GRANTEE/CONTRACTOR TECHNICAL REPORTS

1. This confirms our conversations of 27 Feb 97 and 11 Jul 97. Enclosed are a number of technical reports which were returned to our agency for lack of clear distribution availability statement. This confirms that all reports are unclassified and are "APPROVED FOR PUBLIC RELEASE" with no restrictions.

2. Please contact me if you require additional information. My e-mail is silverr@onr.navy.mil and my phone is (206) 625-3196.


ROBERT J. SILVERMAN

Technical Progress Report 5/1/95 – 7/31/95
Construction of a Connectionist Network Supercomputer
University of California, Berkeley
ONR URI Grant No. N00014-92-J-1617

1 Abstract

This report presents a summary of the technical status for the period 5/1/95–7/31/95. In this period we passed another major milestone in the development of connectionist network computers—we completed the fabrication, assembly, and testing of the SPERT neurocomputing board. We also continued to make significant progress in system and application software for this system, in porting neural network algorithms, and in the application of analog auditory preprocessors to speech recognition.

2 Technical Status

2.1 The T0 Processor and SPERT board

As reported in our previous status report, we received from Hewlet Packard three fabricated wafers containing the T0, vector microprocessor. Using bare die tests developed by us, we tested the wafers and found 40% of the fabricated die passed the tests. This percentage is very good for a die of the size of T0 (approximately 17mm on each side). Also, as reported in the previous quarter, we are developing a workstation add-on processor board, based on the T0 chip. The physical realization of this design is unique in that the T0 chip is wire-bonded directly to the board, without the usual chip package.

This quarter we installed T0 chips on 11 of the fabricated SPERT boards. Five the boards were also fully populated with all memory and support components. The SPERT boards were tested using an innovative test fixture we developed specifically for the purpose, based on elastomeric connector technology. The completed boards are fully functional. Having these SPERT boards working on the bench verifies our use of chip-on-board technology for a large die. We found very few fabrication problems associated with the attachment of the T0 processor to the board—only a single bad wire-bond out of thousands of total wires.

With these boards, we were able to perform more complete and higher-speed testing of the T0 processor. To date, no functional errors have been found in the processor design. The only known problem in the design is an unusually high quiescent power draw (250mA). The high current draw will not limit the usefulness of the chip and will be corrected in later versions.

The maximum clock speed for the T0 processor has been measured at 47MHz. This is significantly higher than our target design speed of 35MHz. This higher speed can be

attributed to our conservative design verification tools and techniques. The speed of SPERT boards will be limited by the commercial SRAM speed. The boards will be clocked at 40MHz.

Two of the SPERT boards have been installed in workstations for general use by our internal research community. These boards have already enabled research that would have been prohibitly slow on our other computing platforms. Now, speech training runs on small neural networks take on the order of minutes as opposed to hours on workstations.

Based on our success, we have begun the steps to put together approximately 20 more SPERT boards, a portion of which will be distributed to other researchers working in neural networks and speech recognition. Towards this goal, we completed a revision of the SPERT board design and sent the design out to fabrication. We have also placed a large order for SRAM parts and expect to put together the SPERT systems in the next quarter.

2.2 Software

Work continues on Torrent library development and system support.

We have performed a series of experiments to determine the effects of fixed point computation on speech recognition neural network training. The results show that we still have some tuning to do on fixed point scaling in these algorithms.

In the area of high-level software, the major advance of this quarter was the new release, 1.07, of Sather. Also, we have setup an internet homepage to provide information on Sather (URL <http://www.icsi.berkeley.edu/Sather/index.html>). It includes tutorials and a list of available technical reports on Sather. The homepage will be continuously updated as new information becomes available.

We have purchased a Myrinet system for parallel Sather (pSather) platform research and as a testbed for CNS-0.

2.3 Analog VLSI pre-processors.

As described in the previous progress reports, we recently constructed a three-chip sound pre-processing system, using three copies of an enhanced version of our 128-channel spectral-shape auditory pre-processor. We are using this system in pilot experiments in speaker-independent, telephone-quality speech recognition.

We conducted similar experiments in 1994, using a single-chip front-end. The best performance we achieved with this system was approximately 13% error, on a 13-words isolated-word task (1-9, "oh", "zero", "yes", "no"), using a 200-speaker, telephone quality database, and a standard Hidden-Markov-Model-based speech recognition system.

Over the past three months we have run recognition experiments using many different configurations of our three pre-processors. These experiments have been done in conjunction

with a neural-network-based Hidden-Markov-Model speech recognition system.

Currently, our best recognition score using our new system is 4.8% error. While this score still falls short of the best software-based front-ends (1.8% error, using the same database and testing methods), it is a considerable improvement from our earlier scores. More importantly, the error score is in the range of usefulness. Current commercial isolated-word telephone quality digit recognizers achieve 2-5% error performance in the field, handling a vocabulary that is slightly smaller vocabulary (11 words) than our 13-word task.

Our research plan for the next quarter is to continue experiments with different recognition approaches, with the goal of further improving our error performance. In addition, we plan to begin speech recognition experiments with noise-corrupted versions of our database, to explore the robustness properties of our front-ends.

Dissemination was also a part of last quarter: a presentation of our multi-chip system was made at the 2nd UCSD-Caltech Joint Symposium on Neural Computation in La Jolla on June 15th, and a paper describing our system appeared in the proceedings. A version of this paper was also submitted to the NIPS conference. Finally, we have begun a study into possible commercial applications of our analog VLSI pre-processor research, and of related micropower analog circuit research, in speech recognition.

3 Recent Publications

Wawrzynek, J., Asanović, K., Kingsbury, B., Beck, J., Johnson, D., Morgan, N., "SPERT-II: A Vector Microprocessor System and its Application to Large Problems in Backpropagation Training," submitted to Neural Information Processing Systems (NIPS 95).

Wawrzynek, J., Asanović, K., Kingsbury, B., Beck, J., Johnson, D., Morgan, N., "SPERT-II: A Vector Microprocessor System," submitted as invited paper to *IEEE Computer* special issue on Neural Computing.

Lazzaro, J. and Wawrzynek, J., "Silicon Models for Auditory Scene Analysis," 2nd Joint Conference on Neural Computation, pp. 134-141, Institute for Neural Computation, 1995. (also submitted to NIPS 95).